

CHAPTER 6:

INSENSITIVITY

Contents

1. Counterexamples and Philosophical Theories.....	2
2. Insensitivity Accounts—Direct and Indirect.....	4
3. The Attack by Counterexamples on Insensitivity Accounts and Two Lines of Response to this Attack.....	7
4. Terminology: We Will Here Use ‘Sensitivity’ and ‘Insensitivity’ to Refer to the Refined Notions.....	9
5. Why Insensitivity Accounts Seem on the Right Track—Even if There Are Counterexamples that Have Not Yet Been Successfully Handled.....	10
6. Perilously Nearby Counterexamples: Schiffer’s and Williamson’s Enhanced BIV Hypotheses, and Hawthorne’s Non-Player Lottery Case	11
7. My Old Refinement: CE	15
8. ‘Real’ Exceptions to IA’s Generalization: Low-Strength, True Claims to ‘Know’ that Skeptical Hypotheses Are False.....	17
9. Williamson’s Distance Underestimator.....	22
10. Relative Insensitivity	26
11. Degree of Belief Insensitivity.....	28
12. Williamson’s Strange Creatures	29
13. Strength of Belief, Basing, and Williamson’s Strange Creatures.....	32

CHAPTER 6:

INSENSITIVITY

To a first approximation, we can say that S's true belief that P is *sensitive* if and only if this conditional holds:

sensitivity conditional: S would not have believed that P if P had been false;

and that S's true belief that P is *insensitive* if and only if this conditional holds:

insensitivity conditional: S would have believed that P (even) if P had been false.

I believe we have very strong reason to think that this notion of sensitivity, suitably refined, will play a starring role in the solution to some important epistemological problems. (And my interests in this book, of course, are especially in AI-inspired skepticism, and the lottery puzzle.) And I believe these are not just *prima facie* strong reasons that have been undermined or overpowered by the recent attacks on sensitivity accounts (or, as they are perhaps better called, *insensitivity* accounts): I think the *balance of reasons* points strongly toward insensitivity playing this important role in epistemology.¹ I thus think that the task of refining the notion of (in)sensitivity is a very important one, and one of the two main goals of this chapter is to contribute to the project of getting the notion right by discussing its application to some important test cases. However, many philosophers seem to think that the attacks against insensitivity accounts have been successful, leaving those accounts bankrupt. So the other main goal of this chapter, and the one we will pursue first, is to argue that this opinion is wrong.

¹ For other defenses of the use of sensitivity in epistemology, see Black and Murphy (2007), and also Murphy and Black (2012), as well as other essays in Becker and Black (2012a), including the editors' introduction to the volume, Becker and Black (2012b).

1. Counterexamples and Philosophical Theories

As I imagine is common among analytic philosophers of roughly my age, one of the first works of contemporary (as opposed to historical) philosophy I read was Saul Kripke's *Naming and Necessity* (Kripke 1980). And as I imagine may be common among those of us with that philosophically formative experience, I developed an aversion to being tied to philosophical 'theories'. Perhaps the most ominous part was:

Let me state then what the cluster concept theory of names is. (It really is a nice theory. The only defect I think it has is probably common to all philosophical theories. It's wrong. You may suspect me of proposing another theory in its place; but I hope not, because I'm sure it's wrong too if it is a theory.) (1980: 64)

But if not theories, then what? Kripke provided the answer: pictures!:

I think I said the other time that philosophical theories are in danger of being false, and so I wasn't going to present an alternative theory. Have I just done so? Well, in a way; but my characterization has been far less specific than a real set of necessary and sufficient conditions for reference would be. . . . So other conditions must be satisfied in order to make this into a really rigorous theory of reference. I don't know that I'm going to do this because, first, I'm sort of too lazy at the moment; secondly, rather than giving a set of necessary and sufficient conditions which will work for a term like reference, I want to present just a *better picture* than the picture presented by the received views. (1980: 93)

These pictures seemed like wonderful things: much like theories, except that you don't have to work out the details, and they're not subject to counter-examples! I can't have been the only beginning philosopher of those days to file away a thought along the lines of: 'Note to self: *Always* trade in pictures; never theories.'

Of course, it was easy to think there was something unfair about Kripke's dealing with his targets' theory (as compared with his own picture). And he addressed this feeling:

Haven't I been very unfair to the description theory? Here I have stated it very precisely—more precisely, perhaps, than it has been stated by any of its advocates. So then it's easy to refute. Maybe if I tried to state mine with sufficient precision in the form of six or seven or eight theses, it would also turn out that when you examine the theses one by one, they will all be false. That might even be so, but the difference is this. What I think the examples I've given show is not simply that there's some technical error here or some mistake there, but

that the whole picture given by this theory of how reference is determined seems to be wrong from the fundamentals. (1980: 93)

My main complaint against Kripke here is that he does not argue that or explain *why* his counterexamples should be taken to show more than just a potentially fixable problem of detail in the theory he has assigned to his opponents (as he seems to think would be all that would be shown by the imagined counterexamples to his own imagined theory), but instead to show that ‘the whole picture’ of his targets is ‘wrong from the fundamentals’. It seems that one shouldn’t just assert, without explanation, that one’s counterexamples seem to achieve the broader destructive end, when all they directly accomplish is the narrower goal.

Ironically, this became an issue for me a couple of years after encountering *Naming and Necessity*, when, as a graduate student, I went to Kripke's own ‘Nozick-bashing’ lectures (as they were often called) against the treatment of knowledge and skepticism in (Nozick 1981).² I wasn't really into epistemology yet (this is what got me in!), and didn't know Nozick's views on knowledge and skepticism. So, I was being introduced to them by means of a killer attack—mainly by means of counterexamples—against them. And *still* I somehow walked out thinking there was something important and right in Nozick’s approach.³ That’s where and how my serious thinking about AI-like skeptical arguments began. And I soon came to realize that what I thought was important and right in Nozick’s treatment was what I called the Subjunctive Conditionals Account (SCA) of the plausibility of AI’s first premise (see esp. sections 5-8 of SSP).

Given how I came to adopt SCA, a pressing question from the very beginning for me was: Well, why think that Nozick was on a good track? Why not take the counterexamples to show that his whole picture is ‘wrong from the fundamentals’?

And, indeed, should not some such responsibility lie with the advocates of philosophical views—whether theorists or more impressionistic painters of

² A paper based on these lectures is now published as ‘Nozick on Knowledge’ in (Kripke 2011), but a couple of my favorite parts, that I can at least seem to remember from the lectures themselves, seem to be missing.

³ Sherrilyn Roush reports a similar experience of immediate attraction to Nozick’s theory of knowledge upon first encountering it (Roush 2005: vii). However, she was reading Nozick himself. I may have her beaten in having my experience in response to an unrelenting attack upon him. (But it should be noted that she reports that she had no training in analytic philosophy at the time in question.)

pictures? I complained above about Kripke, as a wielder of counterexamples, not saying why we should think his examples show a fundamental problem with his targets' views. But if you are promoting a theory, or even something more picturesque, and some counterexamples are causing troubles for your view, shouldn't you give some reason, if you can, for thinking that you are on the right track nonetheless, for thinking that, despite the counterexamples, your *picture* is or might be still a good one?

So in what follows in this chapter, we will look at what insensitivity accounts are, how counterexamples cause troubles for them, why we should think they are on the right track anyway, and how to refine them to handle some of the more troublesome counterexamples.

2. Insensitivity Accounts—Direct and Indirect

So, first: What are insensitivity accounts? What are they accounts of? Well, a variety of cases elicit from us a strong and surprising intuitive pull toward saying that the subjects of the case don't know the propositions in question. Thus, in the relevant cases from recent epistemological literature, there is a strong pull toward saying that I do not know that

E1. I've lost the lottery

E2. My newspaper isn't mistaken about whether the Cubs won yesterday

E3. Those animals are not just cleverly painted mules

E4. I'm not a BIV

Of course, there are many propositions which I intuitively seem not to know. What's surprising about the above? Well, each of the above E's can be paired with another proposition, which we'll in each case label 'O,' about which there are strong intuitions to the effect that a) I do know that O and, b) If I don't know that E, then I don't know that O. Consider these Os, which can each be paired with the similarly numbered E, above:

- O1. I won't be able to repay my loan by the end of the year
- O2. The Cubs won yesterday
- O3. Those animals are zebras
- O4. I have hands

In the case of E2/O2, we suppose that my only source of information about the result of the game is my newspaper, which didn't carry a story about the game, but just listed the score under 'Yesterday's Results.' Intuitively, if the newspaper is a normally reliable one, and, of course, if the Cubs did in fact win, it seems that I know that they won. Yet, in the imagined circumstances (my newspaper is my only source of information about this game), this conditional also seems intuitively correct: If I don't know that my paper isn't mistaken about whether the Cubs won yesterday, then I don't know that they won — if I don't know that E2, then I don't know that O2. These two fairly strong intuitions, if correct, would seem to point to the conclusion that I know that E2. That's why it's surprising that there's such a strong intuitive pull toward saying that I don't know E2. Similar points would apply to our other case pairs.

Of course, the strength of the intuitive pull toward saying that I don't know that E is diminished when it is considered in conjunction with the two other intuitions in play here. In fact, *each* of these three intuitions, because they together form an inconsistent triad of intuitions, is weakened when it is put into contact with the other two. It must be acknowledged that the intuitive phenomena we are seeking to explain here are a bit delicate. But the idea is that each of these three judgments (that I know that O; that I don't know that E; and that I don't know that O if I don't know that E) is individually plausible—it is intuitively plausible when considered on its own, without worrying about how it might conflict with the others. And we will here be concerned with explaining the intuition that one doesn't know that E. So, we're now in a position to ask with a bit of urgency: What accounts for this intuitive pull toward saying that the likes of E1-E4 are not known?

That's what insensitivity is supposed to explain. Given the natural and familiar understandings of the relevant cases, E1-E4 seem to be insensitive beliefs, while O1-O4 seem to be sensitive. (Thus, to continue using the E2/O2 pair, *If the Cubs had not won, I would not have believed that they had won* seems true, while *If my paper had been mistaken about yesterday's game, I would not have believed it*

wasn't mistaken does not seem true: since the paper is my only source for the result, it seems that if my paper had been mistaken, I'd have believed as strongly as I in fact do that it wasn't.) Insensitivity explanations appeal to this insensitivity of beliefs E1-E4 to explain why they seem not to constitute knowledge.

The *direct* way to do this is to follow Fred Dretske (1971) and Nozick (1981) in supposing that sensitivity is a necessary condition for knowledge (or at least that insensitivity is sufficient for a lack of knowledge). If our concept of knowledge were simply that of true, sensitive belief, it would be no surprise that we tend to judge that insensitive beliefs are not knowledge. And, of course, that point will hold also for more complicated theories of knowledge, so long as they make insensitivity sufficient for non-knowledge.

I also appeal to the insensitivity of E1-E4 to explain why such beliefs can seem not to be pieces of knowledge (see SSP and Chapter 5), but, as we've seen, I do so in a very different way. Mine is an *indirect sensitivity* account — one that appeals to the insensitivity of E1-E4 in explaining why they seem not to be knowledge, but does *not* do so by building a sensitivity condition (or anything like a sensitivity condition) into the very concept of knowledge.

Insensitivity accounts seem to depend on some claim to the effect that we have at least a fairly general — though not necessarily exceptionless — tendency to judge that insensitive beliefs are not knowledge. Without some such assumption, the insensitivity of E1-E4 would not do the explanatory work assigned to it. So both types of account utilize what in SSP I called the 'Subjunctive Conditionals Account' ('SCA'), but that we will here call the 'Insensitivity Account' (IA) — by which I mean that in the relevant cases, they explain why S seems not to know that p by means of the following two claims:

IA

1. S's belief that p is insensitive, and
2. We have some at least fairly general — though perhaps not exceptionless— tendency to judge that insensitive beliefs are not knowledge

Where direct and indirect insensitivity accounts diverge is in their further account of why (2) holds. Direct insensitivity accounts hold that this is so because:

- a. Sensitivity is a necessary condition for knowledge

Indirect insensitivity accounts, then, utilize IA, but have some explanation other than the one based on (a) for why (2) holds.

3. The Attack by Counterexamples on Insensitivity Accounts and Two Lines of Response to this Attack

As I reported at the opening of this paper, insensitivity accounts have fallen out of favor, at least in some regions of philosophy. Insofar as this has been based on reasons that have been given for rejecting such accounts, it has been almost entirely due to counterexamples to (a) that have been put forward — cases in which beliefs that are insensitive (at least according to current understandings of ‘insensitivity’) intuitively seem to be, and are put forward by those issuing the attack as being, cases of knowledge nonetheless. Because these are advanced as counterexamples to (a), they are in the first instance attacks on *direct* sensitivity accounts. However, these examples are typically assumed to make trouble even for my *indirect* sensitivity account, because they are exceptions to (2), the generalization utilized by both types of sensitivity account. Now, the astute reader will have noticed that (2) is perfectly compatible with there being exceptions to the tendency it posits. Indeed, I’ve formulated it so that it positively anticipates such exceptions. A main focus of this chapter will be discerning how damaging such exceptions are or would be to insensitivity accounts. In the meantime, it certainly seems that it would be better for explanations to be made by appeal to perfectly pristine generalizations that had absolutely no exceptions, so these exceptions are *prima facie* problematic even for indirect insensitivity accounts.

Sensitivity theorists have responded in two ways to these attacks. First, we have proposed modifications to our accounts of what (in)sensitivity is, by which refinements, it is hoped, the counterexamples can be evaded. It used to be uncontroversial that these modifications had not yet succeeded in dealing with all the proposed counterexamples. At least I took that to be uncontroversial: In a previous defense of insensitivity accounts (DeRose 2004b), after explaining how some proposed refinements to these accounts handle some of the troublesome examples, I admitted: ‘Still, no sensitivity theorist, to my knowledge, has even pretended that all the cases have been successfully dealt with’ (2004b: 25).

However, the state of the debate has changed since then, and that could no longer be said. Troy Cross (2010) has since presented and defended a refined account that he thinks ‘dodges every arrow slung against sensitivity’ (Cross 2010: 40). Now, I should point out that Cross in the end moves to an account which, though it is fairly thought of as a *descendant* of sensitivity accounts, probably should not itself be classified as an insensitivity account, as it abandons appeal to subjunctive conditionals, instead using a formulation built entirely on the notion of explanation, and it is this descendant that he is directly writing about in the above quotation. However, on his way to his favored account, Cross considers various alternatives which clearly should be classified as refined insensitivity accounts and which he argues do not face any clearly problematic counter-examples. About one such account that he labels ‘Relative Sensitivity’ (RS), that we will discuss later, Cross writes: ‘Some readers may stop at this stage, content that the explanatory power of sensitivity has survived its subjection to the epistemologist’s extraordinary counterexamplifying engine’ (2010: 49). Other readers, however, might not be happy with some of RS’s results, Cross fears, so though he thinks RS itself doesn’t face any clearly problematic counterexamples, he goes on to consider some improvements to RS that remain within the sensitivity camp, before moving beyond sensitivity to his own favored, explanation-based account. In case you’re a little worried about what these refined accounts end up looking like (just how ‘refined’ they are, once they’re put through the ringer of multiple rounds of the counterexample game), be assured that Cross is very sensitive to the danger of these accounts getting complicated in ways that make them collapse ‘under the weight of their own caveats and qualifiers’ — which he thinks was the fate of some accounts of knowledge that emerged in the post-Gettier discussion (Cross 2010: 40, n. 22). Cross’s goal is not just to arrive at an account that accommodates the proposed counter-examples, but to do so ‘while leaving some non-baroque remainder that retains at least as much intuitive appeal as sensitivity itself’ (2010: 40), and the accounts he finds successful reflect that goal.

Second, but largely ignored, I have argued that we have good reasons for accepting that the correct explanations for the relevant phenomena are to be found in the direction pointed to by sensitivity accounts even though (or even if) there are exceptions to (2) on our best current accounts of insensitivity. Indeed, I’d say that even in the face of some remaining counterexamples, we would have good reasons for taking the explanations provided by our current understanding of sensitivity *already* to be good ones — though they could be improved by more exact notions of what insensitivity is.

The character of the project of refining insensitivity accounts (and thereby executing the first line of response described above) is greatly affected by whether the second line of response is successful. Without the second line of response, the project of refinement may appear to be very defensive: Insensitivity theorists seem to be on the ropes, with the viability of their approach waiting on hold until they can come up with the refinements needed to avoid the attack they are under. Things look different if the second line of response is successful. Once that second case has been made, we can see that we already have good explanations, based on the notion of insensitivity, for the phenomena these accounts attempt to explain, and the task of refining the notion becomes the happier one of making already good explanations even better.

4. Terminology: We Will Here Use ‘Sensitivity’ and ‘Insensitivity’ to Refer to the Refined Notions

We face a terminological decision. We could just use ‘sensitive’ to describe beliefs about which the sensitivity conditional holds and use ‘insensitive’ for beliefs about which the insensitivity conditional holds. But here we will instead use the terms to refer to the properly refined accounts of these notions, to which the relevant conditionals provide only a rough guide. The idea is that there is an important property of some of our beliefs that the obtaining of insensitivity conditionals gives us a first approximation of, and which tends to make us think those beliefs don’t constitute knowledge. Calling a belief ‘insensitive’ is to say it has that perhaps somewhat elusive property (that we are trying to get a more exact grasp of). Saying that a belief is ‘insensitive according to our current understanding of insensitivity,’ and the like, is saying that it satisfies our current attempt to specify the notion. To say that a belief is insensitive on our first approximation for the notion, I’ll say that the insensitivity conditional holds of the belief. Analogous practices will govern my use of ‘sensitive.’

5. Why Insensitivity Accounts Seem on the Right Track—Even if There Are Counterexamples that Have Not Yet Been Successfully Handled

Even if there are still good counterexamples that show that the generalization IA utilizes is not ideally precise on our current understanding of insensitivity, there is good reason to think that IA is on the right track, nonetheless. So I argued when I put forward an insensitivity account in SSP (see section 8). I will repeat the essence of that argument and expand on it a bit here in this section.

First, it's good to consider why IA is an attractive explanation in the first place—beyond the fact that it produces intuitively correct results to the case pairs it sets out to explain (E/O 1–4, that we looked at in section 1). Much of that attraction, as I've claimed (DeRose 1996: 560), comes from the initial plausibility of IA's generalization: Pointing out that someone would have held a certain belief even if it had been false intuitively seems a very good reason for denying that they know what they believe. IA does not just seem to happen to get a bunch of cases right, but also gets them right by means of an explanation that has a good deal of intuitive pull in its own right. And the most prominent and best critics of insensitivity theories admit this. Kripke:

Nozick's third condition has a clear intuitive basis. "Even if p had been false, you'd still have believed it!" sounds like an objection to a knowledge claim. (Kripke 2011: 180)

And Williamson, writing about the insensitivity conditional:

At first sight, that counterfactual looks like a reason for denying that one knows. (Williamson 2000: 147)

Next, as I stressed (but apparently didn't stress vigorously enough) in SSP,⁴ I was using (2) to explain why we seem not to know in various cases, and the generalization needn't be exceptionless to play that explanatory role. The exceptions perhaps show that the generalization can be refined and improved in certain ways, and may even point us in hopeful directions toward finding some such refinements (some of which are no doubt important and will significantly advance

⁴ See the last paragraph of sect. 5 of SSP.

our understanding, and, indeed, some of which we'll consider here), but heaven help us if we have to wait until the generalizations we use in philosophy (or elsewhere) have to be already perfectly Chisholmed and absolutely exceptionless before we can put them to explanatory work!

But, again, why think the insensitivity account is even on the right track? Why think the exceptions reveal only the need for further tinkering, rather than for a completely different account? Without repeating the case variants I discuss (in sect. 8 of SSP), a good reason is that where the account works, it works so impressively well. On top of the point made two paragraphs above that the holding of the insensitivity conditional intuitively seems like a very good reason to deny that a subject knows, there is the perhaps more impressive development that when we take cases like the familiar specifications of the situations in which our current Es are usually placed, and then start imagining the most natural ways of modifying the situation in question so that the subject *does* seem to know the relevant proposition, we will find in an imposingly impressive array of case variants that the very changes needed to make the subject seem to know also render the subject's belief sensitive. As I concluded: 'Again and again, [IA] posits a certain block [the insensitivity of the belief] to our judging that we know, and the changes that would clear the way for our judging that we know also remove this block. This makes it difficult not to believe that [IA] is at least roughly correct' (SSP, sect. 8). Are we to suppose that it's just a *coincidence* that these Es seem not to be pieces of knowledge when they are in their usual settings, where they are insensitive beliefs, but that they no longer give this 'no-knowledge' appearance in the modified situations in which they are sensitive—that the very changes needed to make the appearance of no-knowledge disappear also render the beliefs in question no longer insensitive?

6. Perilously Nearby Counterexamples: Schiffer's and Williamson's Enhanced BIV Hypotheses, and Hawthorne's Non-Player Lottery Case

But not all counterexamples are equal. If a counterexample is in important ways similar to the cases IA claims to explain, these can be especially damaging to IA, and can jeopardize its claim to have provided a good explanation. It is one thing to stand

by IA's explanation for why one seems not to know that one has lost the lottery or that one is not a BIV (in the standard cases of recent epistemology) in the face of IA's generalization seeming to issue an intuitively wrong result when applied to some very different case about whether one knows that the trash one has dropped down the chute in one's high-rise apartment building has made it down all the way to the bottom (Sosa 1999: 145-6) or whether one knows that some ice one has left out in the sun on a hot day has melted (Vogel 1987: 206-8); it's quite another thing if IA's generalization misfires on some variant lottery or BIV case that is 'nearby' to the original in that it represents a very natural way to modify the case so as to change the intuitive verdict as to whether we know the matter in question.⁵ Indeed, I have supported IA by its ability to get nearby cases right, and have claimed its superiority over other accounts because those others flounder on nearby cases that IA gets right. So it would be downright unfair of me to brush aside *nearby* trouble for IA. So I will here respond to very nearby counterexamples that have appeared in prominent places: BIV cases put forward by Stephen Schiffer and Timothy Williamson, and a lottery example advanced by John Hawthorne. After looking at

⁵ Of course, the division I'm working with here between the 'nearby' counterexamples on the answering of which depends the credibility of the insensitivity approach and other counterexamples is quite vague. Still, I hope even this vague guide clarifies, to the extent possible, my vague approach here. Williamson writes: 'Naturally, individual examples do not refute the hypothesis that most ordinary cases conform to (4), or even to (1). DeRose prudently avoids advancing such principles as exceptionless generalizations; context-dependence is an unruly phenomenon. Nevertheless, he does not dismiss recalcitrant cases as statistically insignificant; he accepts the responsibility to explain them, as his willingness to replace (1) by something like (4) shows' (Williamson 2000: 161). Well, much depends on what exactly is meant here by 'accepts the responsibility to explain.' We should endeavor to explain any of the exceptions — either by modifying our account so that they are no longer exceptions, or explaining, if we can, why such an exception arises, even while our generalization provides explanations of the phenomena it seeks to explain. Doing so improves our account. But I don't accept responsibility for handling the not-so-nearby examples in the sense that I think the insensitivity account is not credible until those handlings have been produced. I'm not sure what Williamson's own stance is here. Toward the beginning of the passage quoted above, he seems to be agreeing that some unhandled examples don't show the insensitivity approach to be wrong, but he does seem in the end to reject the approach on the basis of some cases (the most important of which we will discuss below in sections 9–13) that he thinks the approach cannot handle, but that I would classify as not-so-nearby to the skepticism and lottery applications of the approach. I *think* he does so because it appears to him, not just that the current accounts of insensitivity produce wrong results in the cases, but also that there is little or no hope for such approaches to ever get such cases right. (By contrast, one could be in a position in which one can't yet see how to handle a case, yet one thinks it's quite hopeful that there is a way to do so that one has not yet seen or devised.) That's how I'm inclined to interpret the force of what immediately follows the above quotation: 'It is quite unclear how to explain the counterexamples to (4) within a counterfactual framework. . .' (2000: 161).

these cases in this section, I will, in the next section, explain how a modification I had already suggested in SSP allows IA to get these cases right.

I have claimed that we can account for why it seems to us that we don't know that we're not BIVs by appeal to the insensitivity of one's belief that one is not a BIV: this is a belief one would hold even if it were false (even if one were a BIV). Schiffer and Williamson present variants of the BIV hypothesis, modified so that it will no longer seem to us that we don't know that the hypothesis is false. So, Schiffer points out that while I don't seem to know that I'm not a BIV, I do seem to know that I'm not what Schiffer calls a 'BIV'' or what we'll call a 'BIVnas', which, Schiffer explains, is 'exactly like a BIV except that it lacks auditory sensations' (Schiffer 1996: 331).⁶ Following Schiffer, Williamson points out that, while seeming to himself to be sitting in front of his computer screen in his office, he does seem to know that he is not what we'll call a BIVmc: a BIV who seems to himself to be climbing a mountain (Williamson 2000a: 158). However, at least while we use the simple account of sensitivity we've been working with so far, one's beliefs that one is not a BIVnas and that one is not a BIVmc, like the belief that one is not a BIV, are insensitive: if you were a BIVnas, or if you were a BIVmc, you would still have believed you were no such thing. So IA, when it utilizes the initial account of insensitivity governed totally by the insensitivity conditional, predicts that one should seem not to know that one is not a BIVnas or a BIVmc. But as Schiffer and Williamson correctly point out, in these cases, we don't get the same appearance of ignorance that we get in the simple BIV case: we do seem to know that these new hypotheses don't obtain. Writing of what we're here calling the insensitivity account, Schiffer concludes: 'Further

⁶ In the fifth paragraph of *Meditation Two*, Descartes uses the Latin term 'sentire' (to sense), in such a way that 'sensation' cannot occur without a body. When, for instance, he is wondering which characteristics he can still be said to have under the supposition that he is the victim of an all-powerful deceiver, and so has no body, he rejects the suggestion that he is something that has sensations: 'Sensation? This cannot happen apart from a body; and in sleep I have seemed to have sensations that I have since realized never happened' (Descartes 1971: 69). There seems to be a good use of the English word 'sensation' that works like this, but Schiffer here cannot be using 'sensations' in this way. He must be using the term as Descartes uses the corresponding Latin later, in the ninth paragraph of *Meditation Two*, where, taking stock of what he is, still under the suppositions that he is the victim of the deceiver or is asleep, he writes: 'Finally, it is I who have sensations, or who experience corporeal objects as it were by the senses. Thus, I am now seeing light, hearing a noise, feeling heat. These objects are unreal, for I am asleep; but at least I seem to see, to hear, to be warmed. This cannot be unreal; and this is what is properly called my sensation; further, sensation, precisely so regarded, is nothing by an act of consciousness' (Descartes 1971: 71). We will follow Schiffer in using 'sensation' in this second way.

refinement is needed, and it's not clear how it should go' (1996: 331). But we'll see that the refinement I had already suggested in SSP (Schiffer's main target) handles this case. Williamson realizes that such a refinement to the notion of insensitivity allows IA to get such enhanced BIV cases right, and advances such a refinement, based on the one I had proposed, that handles these cases. He subsequently rejects even such a refined account on the basis of other examples, the most important of which we'll consider in the closing sections of this paper.

But let's first get our last counterexample out. I have also claimed that insensitivity accounts for why it seems that one doesn't know that one has lost the lottery in a lottery situation as it is standardly set up in the epistemological literature (which includes that there are many tickets, that there is only one winner, that one holds only one ticket, and that the winner has not yet been announced). But Hawthorne points out that one typically will seem to know that one hasn't won a lottery in the different situation in which one isn't even playing the lottery—one doesn't even have a ticket. However, even the non-player's belief that she hasn't won seems insensitive, at least on the simple conditional account of insensitivity. As Hawthorne writes: 'For if I had won, I would have owned a ticket, but not having heard the result yet, would think it a loser' (Hawthorne 2004: 11). Like Schiffer and Williamson before him, Hawthorne sees that 'refinement is called for' (2004: 11), and like Schiffer, he doesn't like the chances of a refined account handling his example correctly: Hawthorne goes on to argue that the refinement he considers on behalf of IA won't work, and he thereby rejects the insensitivity approach. Unfortunately, Hawthorne tries out the wrong refinement, and doesn't consider the refinement that I had suggested in SSP that we are about to see handles all three of these counterexamples. (Nor did he consider the closely related proposed refinement which had been suggested by Williamson and which also gets Hawthorne's case right.⁷)

⁷ At (2000: 159), Williamson proposes this refined sensitivity account of knowledge:

(4) Necessarily, if S knows p then, for some proposition q : q entails p , S sensitively believes q , and $\sim p$ does not explain how S could falsely believe q .

This yields the intuitively correct verdict that the non-player in Hawthorne's case knows that she hasn't won the lottery, where q is the proposition that she doesn't have a ticket.

7. My Old Refinement: CE

While some more recent proposals can also handle the above examples, and while Williamson's version of the kind of refinement I've used is in a way nicer than mine,⁸ it is perhaps worthwhile to see that the refinement I had proposed back in SSP gets these cases right, so that old refinement is what I'll apply to the cases here.

Responding to a couple of quite different 'counterexamples' to IA's generalization in SSP, I suggested this refinement:

[IA's] generalization was this: We tend to judge that S doesn't know that P when we think that . . . S would believe that P even if P were false. The limitation of [IA's] generalization that's suggested by these cases is this: We *don't* so judge ourselves ignorant of P where not-P implies something we take ourselves to know to be false, without providing an explanation of how we came to falsely believe this thing we think we know.⁹

Following that formulation very closely yields this refined (conditional-plus-explanation) generalization:

(CE) We tend to judge that S doesn't know that P when

1. We think that S would have believed that P even if P had been false, AND

2. There is no A such that:

a. We take S to know that A is false,

b. Not-P implies A, and

⁸ I like the positiveness of Williamson's way of doing things, looking for a positive, sensitive basis for the subject's belief, rather than my way of seeing if the denial of the subject's belief leads to what the subject can, in some approved way, see to be trouble. We will look at Williamson's proposed refinement in section 9, below.

⁹ This is from the third paragraph of sect. 7 of SSP. The complete version of the first sentence of this quotation, including the material elided above, is: 'SCA's generalization was this: We tend to judge that S doesn't know that P when we think that S's belief that P is insensitive (when we think that S would believe P even if P were false).' I elided the bit about insensitivity for the sake of clarity, since I am here using 'insensitivity' in a different way from how I used it in SSP. There I used the term to cover any true belief that P that satisfied the simple conditional test for insensitivity. Here, we're using the term to denote beliefs that have a property that the simple conditional test provides a first approximation for, but that we hope to provide a more exact analysis of. I have also changed 'SCA' to 'IA'.

c. Not-P fails to explain how S came to falsely believe that not-A

Though I devised (CE) to handle quite different examples, it works on these new cases as well.¹⁰ It distinguishes the simple BIV hypothesis from the BIVnas and BIVmc hypotheses. Consider my beliefs in the non-obtaining of these three hypotheses; that is, consider the cases where P is *I'm not a BIV*, *I'm not a BIVnas*, and *I'm not a BIVmc*. While, as we've already noted, all three of these beliefs satisfy the insensitivity conditional which constitutes the first clause of our account, for each of the latter two beliefs there is an A that satisfies all of the conditions (a)–(c), so these beliefs don't satisfy clause 2. In the case of *I'm not a BIVnas*, that A is *I'm not having auditory sensations*: In the relevant situation, (a) I do take myself to know that I am having auditory sensations; (b) that I am a BIVnas implies that I'm not having auditory sensations; and (c) that I am a BIVnas does fail to explain how I might have come to falsely believe that I'm having auditory sensations. For similar reasons, my belief that I'm not a BIVmc doesn't fall under our refined generalization because (a)–(c) are all satisfied for it where A is *I appear to myself to be climbing a mountain*: (a) I do seem to know that I don't appear to myself to be climbing a mountain (like Williamson, I seem to know that I appear to myself instead to be sitting in front of a computer screen); (b) that I am a BIVmc does imply that I appear to myself to be climbing a mountain; and (c) that I am a BIVmc does fail to explain how I might have come to falsely believe that I don't appear to myself to be climbing a mountain. By contrast, my simpler belief that I'm not a BIV doesn't appear to get rescued by any A so as to fall outside of this generalization's purview. There are candidates for A that satisfy (a)–(b) — *I don't have hands*, for instance: (a) I do seem to know that I do have hands; and (b) that I am a BIV does imply that I don't have hands. Here, however, (c) fails, because, famously, the hypothesis that I am a BIV does seem quite capable of explaining how I might have come to falsely believe that I have hands.

In Hawthorne's example of the lottery non-player, P is *I haven't won the lottery*. The non-player does seem to know this of himself. But that's fine by our refined account, because for the non-player, there is an A that satisfies (a)–(c) and thereby allows the non-player's belief that he hasn't won to escape the generalization, namely, *I have a ticket for this lottery*: (a) the non-player does seem to know that this

¹⁰ See also (Murphy and Black 2012), which applies this kind of account—though in the form of an account of what knowledge is, rather than as an account of when we tend to judge that there is no knowledge—to a variety of other cases.

A is false — he does seem to know that he doesn't have a ticket; (b) that he has won the lottery implies that he has a ticket; and (c) that he has won the lottery does fail to explain how he came to falsely believe that he has a ticket. Note that if not-P is beefed up to *My friends have secretly bought me a ticket to the lottery, and, because that ticket's number has been drawn, I have won the lottery*, then clause (c) is satisfied on natural ways of construing the resulting situation, but that's alright, because in the same scenarios that (c) is satisfied, one's belief that that beefed-up not-P is false does seem not to constitute knowledge — that beefed-up not-P seems to be an intuitively successful skeptical hypothesis even for the non-player.

8. 'Real' Exceptions to IA's Generalization: Low-Strength, True Claims to 'Know' that Skeptical Hypotheses Are False

As I've indicated, I think there is an especially pressing call to fine-tune IA's generalization so as to handle cases like the ones we have just been discussing. Now we turn to the other extreme: cases in response to which I think it would be a bad idea even to try to modify our account of what insensitivity is so as to predict the intuitively correct verdict.¹¹

I have always¹² thought that there could be, should be, and are, 'real' exceptions to IA's generalization: cases about which our intuition will be that a subject knows, even though their belief in question is (and is from the point of view of those doing the intuiting) not only insensitive on our current understanding of 'insensitivity,' but would retain that classification even on the best, ideally refined, construals of 'insensitivity.'

¹¹ In a broad sense, of course, it is good to augment our account to handle these cases. But in these cases, I think this should be done by explaining why the beliefs in question seem to constitute knowledge despite the fact that they are insensitive, rather than to modify our account of insensitivity so that it rules that these beliefs are not insensitive.

¹² Well, since at least 1990, when I wrote my dissertation, but I think the conviction goes a year or two further back than that.

That such cases could and should occur is based on my account of why IA's generalization holds—why we tend to think that insensitive beliefs don't constitute knowledge. As I explained above in section 1, as an indirect insensitivity theorist, I don't think the generalization holds because a sensitivity requirement is part of our concept of knowledge, but for some other reason. On my account, the generalization holds due to the operation (described in sect. 12 of SSP) of a mechanism for raising the standards for what counts as knowledge.¹³ Without repeating here the whole basis for how this works, nor all the details, in brief, this 'Rule of Sensitivity' (as worded in the first paragraph of sect. 12 of SSP) states that

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge. Where the P involved is to the effect that a skeptical hypothesis does not obtain, then this rule dictates that the standards will be raised to a quite high level, for, as we've seen, one must be in a stronger epistemic position with respect to a proposition stating that a skeptical hypothesis is false—relative to other, more ordinary, propositions—before a belief in such a proposition can be sensitive.

The reason that insensitive beliefs tend to appear not to be instances of knowledge, and the reason that sensitivity can in general *appear to be* a requirement for knowledge, then, is that this 'Rule of Sensitivity' tends to operate to make claims that 'S knows that P' turn out false (and claims that 'S does not know that P' turn out true) when S's belief that P is insensitive. For instance, on my view, we *do* know that we're not BIVs by the epistemic standards that govern most conversations. The problem is that an attempt to claim this 'knowledge' will, by the Rule of Sensitivity, tend to put into place precisely the exceedingly high standards for knowledge that we don't meet. Since claims to the effect that insensitive beliefs are 'known,' for this reason, tend to be false, and claims that insensitive beliefs are *not* 'known' tend to be true, it's no wonder that insensitivity will appear to be a knowledge-killer.

But note all those 'tend to's. This mechanism for raising the standards for knowledge is just one among many conversational forces affecting epistemic standards. So, it can come into conflict with other conversational forces that oppose it, where the operation of the Rule of Sensitivity is pushing the 'conversational score' in one direction (toward higher standards for knowledge) while other forces are

¹³ Here I use 'counts as' in the way described at (DeRose 2009: 187).

pushing for lower standards, or perhaps for keeping the standards where they are. What happens then? Hard to say.¹⁴ Sometimes it might just be somehow unclear what epistemic standards govern such a conversation—whether because there are determinate standards in place but it’s hard to discern what they are, or because it is objectively indeterminate what standards are in place.

But couldn’t, and shouldn’t, there be at least *some* cases where we can discern that the ‘Rule of Sensitivity’ has been over-ridden, or for some reason has failed to operate, and a claim that ‘S knows that P’ is true (or that ‘S doesn’t know that P’ is false), despite the fact that S’s belief that P is insensitive? (Isn’t that how such pragmatic principles work in general?) On my picture, a main function of inserting skeptical hypotheses into a conversation is to raise the standards for knowledge that are governing that conversation — raising them a moderate amount in the case of bringing up a moderately remote skeptical hypothesis, and raising them an extreme amount in the case of bringing up an extremely remote skeptical hypothesis.¹⁵ But might there not be contexts in which it’s fairly clear that a speaker is bringing up the matter of whether a subject does or does not know a skeptical hypothesis to be false, not for the usual reason for bringing up such a thing, but for some other purpose?

At any rate, the best cases of true assertions of ‘S knows that P,’ where S’s belief that P is ‘really’ insensitive, that I’ve been able to come up with are cases in which it’s fairly clear that the reason this is being said is not to assert that S knows that P by even the unusually high standards that such an assertion would tend to put into play, but to (favorably) *compare* S’s standardly good epistemic position with respect to P with that of some poor unfortunate who is not even in that moderately good position that S is in. In one such case I have used (DeRose 1990: 221–3), we can imagine that some people, perhaps due to the influence of strong mind-altering drugs used in powerful combination with a reading Descartes’s *Meditation One*, have come to be very unsure about whether they are victims of Descartes’ deceiving evil

¹⁴ But for some reflections on what happens in the philosophically interesting cases of discussions of skepticism, see (DeRose 2004a), updated and improved at (DeRose 2009: 128–152).

¹⁵ Of course, trying to raise the epistemic standards by an extreme amount is likely to meet with resistance — especially if one’s circumstances don’t make such high standards appropriate. What happens in cases where one meets with such resistance in the form of an interlocutor determined not to let you get away with raising the standards (so much) is precisely the topic of (DeRose 2004a) ≈ (DeRose 2009: 128–52).

genius, and that others, who ingested a bit more of the drugs, have come to believe, not just that they might well be, but that they actually are, such victims of the evil genius. These poor souls may all be gathered together, perhaps in some support group at a psychiatrist's office. They are each in a *deplorably* bad epistemic position with respect to their not being a victim of Descartes's evil genius: they don't even believe it (and in the extreme cases, believe its opposite)! By contrast, on my account, normal people (even without the aid of any heroic anti-skeptical reasoning by philosophers) naturally tend to be in a very strong epistemic position with respect to the nonobtaining of such skeptical hypotheses. Now, if someone observing this meeting is unsure whether a particular person is such a patient, they might point to the person and ask the psychiatrist: 'Is *he* one of those poor patients who think they are or may be victims of the evil genius? He looks pretty normal.' When the psychiatrist responds, 'Oh, no. He's not a patient; he's my assistant. Don't worry: He knows that he's not the victim of an evil genius,' the claim that this assistant 'knows' that he's not the victim of an evil genius certainly doesn't seem to have anything of the flavor such a claim would have when it's made to report a philosopher's having triumphantly reached the conclusion of some very fancy (and likely, in the end, unsuccessful) anti-skeptical reasoning (perhaps of a Cartesian or Putnamian variety). The psychiatrist's attribution doesn't seem to be ascribing to the assistant the kind of supremely exalted knowledge that talk of evil geniuses tends to put into play, but rather to be the much more restrained — but in this context, conversationally very appropriate — and true claim that, as compared with the gathered patients, her assistant is in the standardly strong epistemic position with respect to his not being the victim of an evil genius.¹⁶ A similarly humble interpretation of a patient's own self-attribution of 'knowledge' seems called for as

¹⁶ As observed in n. 23 of Chapter 3, John Koetha is among those who missed my cases in which one can truthfully claim that someone 'knows' a skeptical hypothesis to be false (even though the subject's belief that H is false is insensitive), and misattributes to me the view that we can never make such true claims. What's fascinating is that when Koetha produces an example that he thinks shows that one can say that one knows what one's real condition is, as opposed to some skeptical hypothesis, he independently reaches for a case just like the ones I had used and am about to present here in that it utilizes the device of comparing a normal person's condition with that of pitiable folks who are poorly positioned to know that they're not in a skeptical scenario. Koetha writes: 'Contextualism is supposed to account for the power of skeptical arguments while reassuring us that our common-sense claims to knowledge are true; but this reassurance is hollow, for I can never say or think, for instance, that unlike the poor souls in *The Matrix*, I at least know what my real condition is' (Koethe 2005: 76). As I've explained, on good versions of contextualism, one can truthfully say or think such things.

we imagine one of the patients gradually recovering from her condition (as the drugs slowly work their way out of her system), and reporting at the support meetings, with the goal of marking her progress (thereby comparing her condition with her own past condition, and perhaps also that of other patients), first, 'I *think* I'm not being deceived,' then 'I'm now pretty sure I'm not being deceived,' and then, finally, 'I *know* I'm not being deceived. What an absurd idea! Thank God I'm normal again!'

It's worth also considering another case I've used (DeRose 1990: 223–4), since, riffing on a well-known case from the epistemological literature (Dretske 1970: 1015–16), it involves a more moderate 'skeptical hypothesis.' So, suppose a group of friends is meeting by the zebra cage at the local zoo, and as two of the friends, Nico and Lou, who are a bit late, approach from a distance, one who is already at the meeting place, Andy, devises a practical joke to play on them: the rest of the friends will try to get Nico and Lou to falsely believe that the zebras are cleverly painted mules. When Nico and Lou finally arrive, the friends around them start saying things like, 'What a great paint job! I would have never guessed that they're just painted mules,' and, 'When I heard of this painted mules scam, I never thought the zoo would be able to get away with it so easily.' After Andy walks away to buy several stuffed zebras from a nearby concession stand, Nico catches someone snickering, and has to be let in on the joke, but Lou is completely taken in by it. So, when Andy returns, one of the friends pulls him aside to apprise him of the situation, explaining: 'Nico knows that they're not painted mules. Sorry, we had to tell her. But Lou totally bought it!' Here Nico, like all her friends, has not performed any of the investigations that would make her belief that the animals aren't cleverly painted mules sensitive, and as a result, her belief is ('really') insensitive. Yet here it seems, at least to me, that 'knowledge' that the animals are not painted mules is truthfully attributed to her. Given the conversational circumstances, the assertion that 'Nico knows that they're not painted mules' does not have the effect of raising the epistemic standards to the unusually high level at which Nico's belief in that proposition must be sensitive to count as 'knowledge.' The point our speaker fairly clearly means to be making is not that Nico has the kind of *fairly* exalted knowledge of the fact that the animals are not painted mules that some special investigation (perhaps, as has been suggested (Stine 1976: 252), one involving paint remover) or special knowledge (as in the case of my 'boastful zoologist' (SSP, sect. 3)) would allow her to have, but just that she has the more humble (and strangely difficult to

claim, in *most* conversational circumstances) ‘knowledge’ of such facts that most trusting zoo-goers at reputable zoos have — but that Lou here lacks.

The relevant intuitions about these cases may be delicate. I myself am pretty firm in my judgments that the knowledge attributions made in these cases — the psychiatrist’s claim that her assistant ‘knows’ he’s not the victim of an evil genius, the recovering patient’s claim that she ‘knows’ that of herself, and Andy’s friend’s claim that Nico ‘knows’ the animals aren’t cleverly painted mules — are true, so I certainly don’t think that all attributions of ‘knowledge’ to insensitive believers are false. Over the closely related but distinct object-level matter of whether these characters know the propositions in question (as opposed to whether the attributions of ‘knowledge’ made within the story are true), intuitions may vary considerably. To the extent that intuitions about these cases are that the characters do know, I would take these to be ‘real’ exceptions to IA’s generalization, but exceptions that do not damage the explanatory power of IA in the cases to which it is applied—at least when IA is held in the form that I hold it, on which such exceptions are not at all unexpected.

9. Williamson’s Distance Underestimator

We now turn to a couple of middle cases (the second of which is a variant of the first): Unlike the cases we just discussed in the previous section, these are cases that we should try to accommodate if possible by means of refining our notion of insensitivity; but unlike the cases discussed in sections 5–6, the credibility of the insensitivity approach does not, in my judgment, await our producing the needed refinements. These very interesting cases, due to Williamson, are opportunities to make our notion of insensitivity more precise — and also, I think, to start to think through important issues concerning how our ability to believe things to different degrees is connected to our ability to know things.

Williamson presents this very interesting case:

I tend to slightly underestimate the distances I see. When I see a distance of twenty-one metres I judge it to be less than twenty metres, although when I see a distance of twenty-three metres I do not judge it to be less than twenty metres. This may mean that when I see

a distance of nineteen metres and correctly judge it to be less than twenty metres, I do not know it to be less than twenty metres. It surely does not mean that when I see a distance of one metre and correctly judge it to be less than twenty metres, I do not know it to be less than twenty metres. (Williamson 2000a: 159)

I certainly agree with Williamson's judgment that his slight underestimator (as one would naturally imagine him, given the description) knows that the one meter distance he is seeing is less than twenty meters. As Williamson points out, the insensitivity conditional may hold for this underestimator's belief and still he'll seem to know that the one meter distance is less than twenty meters, and he will cause trouble for several refined accounts of insensitivity, too.

In particular, this example is a problem for the account Williamson offered me (based on my proposed modification) by which I could handle cases like his BIVmc case. Williamson suggests:

(4) Necessarily, if S knows p then, for some proposition q : q entails p , S sensitively believes q , and $\sim p$ does not explain how S could falsely believe q . (Williamson 2000a: 159)

Here, Williamson is using 'sensitively' so that 'S sensitively believes q ' simply means that S's belief that q satisfies the initial account of sensitivity, based on the simple conditional test. Also, Williamson formulates the accounts he considers, as above, as necessary conditions for knowledge, rather than, as I have been doing here, as circumstances under which we will tend to judge that subjects don't know. He is addressing direct and indirect insensitivity theorists at the same time, since the same cases tend to be problematic for both sorts of insensitivity accounts. So, though I won't accept any such condition for what knowledge is (as Williamson realizes), the idea is that for me to accept such a refinement is to accept that when a subject's belief that P fails to satisfy (4), *we will tend to judge that* the subject doesn't know that P . So, spelling out Williamson's use of 'sensitively,' and converting his suggestion to an account of when we will tend to judge that subjects don't know things, yields the following refined proposal:

(WP) We tend to judge that S doesn't know that P when there is no Q such that:

1. S believes that Q
2. S would not have believed that Q if Q had been false
3. Q entails P

4. Not-P does not explain how S could falsely believe that Q

(WP) predicts that we will judge that Williamson's underestimator doesn't know that the distance he is seeing is less than twenty meters. For there is no Q that works here where P is *That distance is less than twenty meters*. First, Q can't be P itself here, since, as we've already discussed, S's belief that the distance is less than twenty meters fails to satisfy the sensitivity conditional. And Q can't be a proposition like *That distance is less than three meters*, for though that entails that the distance is less than twenty meters, and though it seems that *That distance is not less than twenty meters* does fail to explain how S could falsely believe that the distance is less than three meters, S's belief that the distance is less than three meters is also insensitive, on our simple initial account of (in)sensitivity—or at least it can be in this case. We can suppose that the circumstances are such that if the distance Williamson's character is seeing had not been less than three meters, he still would have judged that it was less than three meters (because if it hadn't been less than three meters, it would have been only very slightly more than three meters, which would have led our *systematic* distance underestimator astray). So WP predicts that he will not seem to know that the distance is less than twenty meters.¹⁷ But, as I've admitted, this is the wrong result: Williamson's underestimator does seem to know that the distance is less than twenty meters.

¹⁷ Well, that's what (WP) predicts as it stands. Immediately after displaying (4) (on which we've based (WP)), Williamson writes:

Further modifications could be made. We might require that S believes *p* because S believes *q*. We might allow the link between *q* and *p* to be looser than entailment. The discussion below will not depend on these details. (Williamson 2000a:159)

Williamson then rejects this proposal on the basis of his underestimator case (159–61). However, if we do 'allow the link between *q* and *p* to be looser than entailment,' (WP) and (4) may be able to get this case right, for then Q can be something like (phrased in the first person from the point of view of the subject): *That distance appears to me to be less than three meters*. Conditions (1), (2), and (4) of (WP) seem to be satisfied here, so if (3) is loosened up a bit, these accounts may be saved from the case of the underestimator. Of course, this all depends on coming up with a looser version of condition (3) that will work — where working means not just getting this case right, but doing so in a way that doesn't cause trouble elsewhere for the account. And that's very tricky business, it turns out. But if we are envisioning the possibility of loosening up condition (3), it seems premature for Williamson to conclude that such variations cannot save insensitivity accounts from his case. This matter does seem to come down to details.

What about my old refinement? Well, it can initially appear that this case can be handled by (CE), which I'll reproduce here for convenience:

(CE) We tend to judge that S doesn't know that P when

1. We think that S would have believed that P even if P had been false, AND
2. There is no A such that:
 - a. We take S to know that A is false,
 - b. Not-P implies A, and
 - c. Not-P fails to explain how S came to falsely believe that not-A

To get our bearings: The intuition about this case is that the underestimator does know that the distance is less than twenty meters. So (CE) is in trouble if both of 1–2 are satisfied. And 1 is satisfied, so the hopes of this principle are pinned on 2 not being satisfied: on there being some A that does satisfy all of (a)–(c). And it seems that *That distance is not less than three meters* can work as the saving A here. (a): We do take the underestimator to know that it's false that the distance is not less than three meters: Though he's an underestimator, he's only a slight underestimator, and he seems to know that this distance, which is in fact only one meter, is less than three meters. (b): *That distance is not less than twenty meters* does imply *That distance is not less than three meters*. (c): *That distance is not less than twenty meters* does fail to explain how S came to falsely believe that the distance is less than three meters. The problem here is with (a). Though we do take S to know that the distance is less than three meters (and thereby to know that A is false, where A is *That distance is not less than three meters*), that we ascribe knowledge here is itself problematic for an insensitivity account, for S's belief that the distance is less than 3 meters is (or can be) itself insensitive, on the initial account of insensitivity. When we apply (CE) to S's belief that the distance is less than 3 meters, clause 1 is (or can be) again satisfied, and we need a new A to save our account, and though that could be tried (perhaps *That distance is not less than 2 meters* might work as the new saving A), the problem will resurface again, and it seems that this game isn't going to end happily for (CE).

10. Relative Insensitivity

But there are a couple of attractive ways for the insensitivity theorist to handle examples like that of Williamson's underestimator. One important possibility, due to Cross (2010: 47–9),¹⁸ is to utilize the notion of relative (in)sensitivity.

Consider again:

(WP) We tend to judge that S doesn't know that P when there is no Q such that:

1. S believes that Q
2. S would not have believed that Q if Q had been false
3. Q entails P
4. Not-P does not explain how S could falsely believe that Q

On WP, one of the things we're looking for in cases of knowledge that P is that the belief in P itself satisfies the sensitivity conditional (where $P = Q$), or that it has a sensitive basis in the form of some other belief that satisfies the sensitivity conditional, that entails P, and that doesn't get undermined by the supposition that P is false. When we don't find either of those things, we tend to judge that the belief that P isn't knowledge. Williamson's underestimator challenges this account, because the candidates for Q that would satisfy clauses 1, 3, and 4 — like *That distance is less than three meters* — do, or can, themselves fail to satisfy the sensitivity conditional, so clause 2 fails.

But then a solution to this problem becomes apparent. Maybe Q can serve as an adequate basis for P because, though S would have believed that Q even if Q had been false, S would *not* have believed that Q if P had been false: Surely, in the closest worlds in which the distance is not less than twenty meters, S doesn't believe that it

¹⁸ This is one of many ideas for how to modify insensitivity theories that were discussed in the seminar on 'Safety and Sensitivity of Beliefs' that Cross and I co-taught at Yale in the Spring of 2006. My recollection is that relative (in)sensitivity was Cross's idea—though I don't recall how or when it was introduced into the discussion of the seminar. Cross's recollection is that he had a conviction that something in the neighborhood would handle Williamson's underestimator case, and that this vague conviction of his was refined to the notion of relative (in)sensitivity during discussion.

is less than three meters. So, the idea would be to replace (2) in (WP) with something like:

2'. S would not have believed that Q if P had been false

Indeed, once you have the idea of relative sensitivity, you might think it can do a lot of the work itself. Cross proposes a relative sensitivity account of knowledge:

(RS) Necessarily, if S knows p, then S's belief in some ground for p is sensitive to p. (Cross 2010: 49)

Converting this to an account of when we will tend to judge that there is no knowledge, and spelling out 'is sensitive to p,' yields Relative Insensitivity:

(RI) We tend to judge that S doesn't know that P when there is no Q such that:

1. S believes that Q
2. Q is, for S, a ground for P and
3. S would not have believed that Q if P had been false

This has to be fleshed out by an account of what it is for one belief to be a 'ground' for another, but such an account seems extremely promising for handling a variety of examples, including that of Williamson's underestimator. (See Cross 2010 for some applications.) For our purposes, it is worth pointing out that (RI) does well by the BIV and lottery cases we've considered, providing (on some good ways of working out the notion of one belief being the 'ground' of another) an explanation for why we seem not to know in the standard examples, while nicely allowing for knowledge of the relevant propositions in the cases of the BIVnas, BIVmc, and the non-player lottery cases.¹⁹

¹⁹ Where P is *I am not a BIVnas*, a Q that works is *I am having auditory sensations*. For *I am not a BIVmc*, try *I do not seem to myself to be climbing a mountain* in the Q spot. In Hawthorne's non-player lottery case, try *I don't have a ticket*.

11. Degree of Belief Insensitivity

Another way of handling Williamson's example is based on the observation that there is a way in which the underestimator's believing that the distance is less than twenty meters is sensitive to whether that distance is less than twenty meters: Though (we are imagining) he still believes that the distance is less than twenty meters in the closest worlds in which it isn't less than twenty meters, if he's anything close to a normal person, he doesn't at those worlds believe that it is less than twenty meters nearly as confidently, or to the same degree, as he believes that in the actual world, where the distance is only one meter. *Whether* he believes this P isn't sensitive to whether P is true, *but the degree to which he believes* this P certainly is sensitive to whether it's true — again, if he is anything close to a normal person. This suggests we try modifying the insensitivity conditional — 'S would have believed that P (even) if P had been false' — to something like what we might call the 'degree of belief insensitivity conditional': 'S would have believed that P *to the degree that she does* (even) if P had been false.'

This is a very attractive modification because it is right in line with — and, in fact, builds upon — the intuitive basis that insensitivity accounts enjoy in the first place. As we noted back in section 5 along with Kripke and Williamson, pointing out the insensitivity conditional that S would have believed that P even if P had been false intuitively seems a very good reason for denying that S knows that P. But note now, first, that the power of this reason is at least considerably blunted by the reply: 'Yeah, but she wouldn't in that case have believed it to the degree that she does.' Second, the intuitive power of the reason is greater if the degree clause is added to it: 'She would have believed it [to the degree that/every bit as confidently as] she does even if it had been false.' Indeed, the addition of such a clause is *so* intuitive, at least to me, that before I had even encountered any cases where its addition is needed for insensitivity accounts to deliver the right result, and so before I had considered any modifications along these lines, I naturally found myself injecting such a clause into my informal presentations of the insensitivity account. So, for example, here I am in SSP (fifth paragraph of sect. 5), explaining, in terms of insensitivity, why it seems that we don't know that we're not BIVs; the single words that are italicized are so emphasized in the original, but take note of the eight-word phrase in the second-to-last sentence of the below passage, to which I am now adding emphasis:

Upon hearing the [BIV] hypothesis, typically one can't help but projecting oneself into it. How would things seem to me if that situation obtained? Well, pretty much (or sometimes exactly) as they actually seem to me. And, so, what would I believe if such a 'strange' situation obtained? Pretty much (or exactly) what I actually believe. For example, and in particular, if I *were* a BIV, I would believe *every bit as firmly as I actually do* that I *wasn't* one. But if this belief is one I would hold even if it were false, how can I be in a position to tell that, or discern that, or *know* that, it's true?

Applied to Williamson's underestimator, the idea is that, as we'd naturally imagine him, the degree-of-belief insensitivity conditional isn't true of this character — and, in fact, the corresponding degree-of-belief sensitivity conditional is true of him. So it's no problem to a degree-of-belief insensitivity account that we judge that the underestimator does know that the distance is less than twenty meters.

12. Williamson's Strange Creatures

Williamson replies²⁰ that slight underestimators like the character in his example (whom he imagines to be himself) seem to have knowledge even when we imagine them in ways that make the degree-of-belief insensitivity conditional true of them:

A different proposal is to take *degree of belief* into account. The idea is that if the mark had been slightly more than twenty metres above the waterline, I would still have believed that it was less than twenty metres above the waterline, although with less confidence than I believe it when the distance is only one metre. But what if I am not like that? Suppose that once I form a belief in a marginal case, I stick to it; perhaps a macho mechanism causes me to feel an aggressive confidence in it even greater than I feel in non-marginal cases. Regrettable though that may be, when the distance is one metre it does not prevent me from knowing that it is less than twenty metres. Creatures whose beliefs are all or nothing in degree can have such knowledge. (Williamson 2000a: 161)

If Williamson is right about what the intuitions are in these cases (that both the macho version of himself and his all-or-nothing-believing creatures know that the very short distance is less than twenty meters), that would be a good reason not to

²⁰ I suggested the degree of belief modification in comments on a draft of (Williamson 2000a) that I emailed to Williamson. His reply is in the published version of the book.

pursue the degree-of-belief approach we are considering. Perhaps we should in that case opt for the Relative Insensitivity approach? (In sect. 13, below, we'll consider this possibility, and see that it might in the end deliver the same verdict as the degree of belief insensitivity approach.)

But I worry that looking for an account that predicts intuitions of knowledge in such cases would be a mistake, generated by a failure to fully come to grips with what the subjects we are imagining would really be like.²¹ I suspect that our ability to believe things to different degrees, and to adjust our degree of belief to the epistemic situations that confront us, are very deep and important facts about us, and that therefore, we shouldn't form quick judgments about what cognitive achievements would be attained by creatures that are in some ways like us, but are lacking these absolutely crucial features and abilities. Intuitions may tend toward much more negative epistemic appraisals of such creatures upon engaging in even a little bit of more careful thought about what they would be like.

My considered reaction to Williamson's suggestion was that such creatures would be bizarre in such a way and to such a great extent that getting to really know one of them would, at the least, cast quite a bit of doubt on whether they know even that the very short distance is less than twenty meters. Cross puts some fun descriptive meat on the bones in the form of actual cases in which their strangeness is brought to light:

DeRose isn't sure whether such beings could have much knowledge at all, at least knowledge of the 'x is less than y ft.' variety. It is, in fact, difficult to imagine what such beings would be like. Their betting behavior and belief updating would be truly bizarre. Imagine that a two-foot pole slowly grows taller and taller. As the pole grows taller, these creatures continue to bet that it is less than twenty feet tall and will take the very same odds until exactly one moment when the pole is slightly taller than twenty feet. Then, there is a discontinuous

²¹ In what follows, I focus on the case of Williamson's strange creatures whose beliefs are all-or-nothing in degree. Fairly similar points would hold for the case of someone being 'macho' in the way Williamson describes. We perhaps all know people who are, to a much lesser extent, macho in something of the way that Williamson points to here: suffering from an unattractive epistemic defensiveness. But to suppose that someone really isn't one bit more confident that the distance they're seeing is less than twenty meters when it's just one meter than they are when it is slightly over twenty meters is to imagine a bizarre believer indeed — in some ways much like the strange creatures we will discuss. (But since a subject who is 'macho' in the way described only has his extreme over-confidence kick in when cases are marginal, such a subject would be in some ways even stranger than the all-or-nothing believers.)

change in credence and they bet the other way, at some particular odds that remain fixed regardless of how tall the pole grows beyond twenty feet. You don't want such a creature working as your parking attendant. Suppose you tell him the hood of your Bentley, which slopes out of the driver's view, is exactly four feet long. In parking your car, he'll underestimate the distance to the wall, smashing confidently into it. From exactly four feet away from the wall, he will be just as confident that he is more than four feet away as he was when he was a mile away. (If his degrees of belief are .7 and .3, he will, bizarrely, be somewhat wary of smashing into the wall when he is still a mile away from it.) Who knows what to make of such creatures? (Cross 2010: 46)

But would these creatures really appear to fail to know that even a very short (less than one meter) distance is less than twenty meters? I don't see why not. My guess is that we would stop thinking of such creatures as having such knowledge once we got to know one of these guys well enough to truly realize that, though he is, as seems appropriate to the situation, quite confident that the very short distance he is seeing is less than twenty meters, he would be (as we have to imagine him for him to be trouble for our degree-of-belief insensitivity account) *every bit* as confident that it was less than twenty meters if the distance were even much greater, to the point of being slightly more than twenty meters; that the difference between seeing the very short distance that he actually sees and a distance of slightly more than twenty meters makes absolutely no difference at all to the degree to which he believes that the distance is less than twenty meters. I mean, try to imagine spending some time with such a creature. Imagine experiencing the supreme confidence with which he made false distance judgments (which we have to imagine he continues to make, to keep him a slight distance underestimator), and his complete inability to at all moderate his level of confidence in response to such bloopers. Then imagine talking to him as he looked at the very short, one meter distance:

You: So, do you think that distance is less than twenty meters?

Creature: Yes. In fact, I'm supremely confident of that.

You: Remember yesterday, when you said you were supremely confident that a distance was less than twenty meters, and you turned out to be wrong, and also those other times you were so supremely confident about distances but turned out to be wrong? Are you *any* more confident this time about your judgment that the distance is less than twenty meters than you were yesterday, or on all those other occasions?

Creature: No, not one bit more confident. Still, I am *supremely* confident again now.

I don't think this guy knows, nor that he will seem to know to anyone who really gets to know him and becomes familiar with the depths of his severe cognitive problems. To get such a bizarre all-or-nothing-believing creature to seem a knower in cases like this, I have to suppose that he has learned to believe (at all) that a distance is less than twenty meters only when he is also inclined to believe that it is considerably less than that, as well. But then he isn't any more the kind of distance underestimator that makes trouble for our account. And I strongly suspect that would be the intuition of most folks who really got to know him. At the very least, I am completely disinclined to go out of my way to tailor an insensitivity account to make sure it allowed for the appearance of knowledge in this truly bizarre case.

13. Strength of Belief, Basing, and Williamson's Strange Creatures

Let's briefly consider how the example of such a strange creature might be handled by RI, which I'll copy here for convenience:

(RI) We tend to judge that S doesn't know that P when there is no Q such that:

1. S believes that Q
2. Q is, for S, a ground for P and
3. S would not have believed that Q if P had been false

Does RI avoid predicting that the intuition about this case will be that the strange underestimating creature doesn't know?

As I remarked in section 10, RI has to be fleshed out by an account of what it is for one belief to be a 'ground' for another, but the basic idea here (at least in part: there must also be a component about Q providing good enough support for P; one version of that component would require Q to entail P) is that the believer somehow bases her belief that P on her belief that Q.

But when does that happen — under what conditions does S base her belief that P on her belief that Q? I'm very far from being in a position to give a useful complete analysis of that notion, but one thing that has made sense to me to at least

look for in a case of a belief that P being based on a belief that Q is that the believer be sensitive in the degree to which she believes P to the support that, from her point of view, P receives from Q: that she would not have believed P to the degree that she does, were it not for the support that belief receives in her belief system from her belief that Q.²² Or something like that. And the problem, then, with supposing that the strange creature we are discussing has some good basis or ground for his belief that the distance is less than twenty meters is that nothing at all like that can be going on, since belief for him is an all-or-nothing affair. So, for instance, *That distance is less than three meters* can appear to be a promising candidate for Q here; conditions 1 and 3 seem to be met. And if we don't require that Q entail P,²³ another very good candidate for Q can be (phrased from the subject's point of view) *That distance appears to me to be less than three meters*; again, conditions 1 and 3 seem met. But can our creature really be said to be basing his belief that the distance is less than twenty meters on his belief that it is less than three meters, or on his belief that it appears to him to be less than three meters, when he would have believed every bit as strongly that it was less than twenty meters if it had been twenty meters, and it had therefore produced in him no belief nor appearance (nor belief in an appearance) that it was less than three meters? The beliefs that the distance is or appears to him to be less than three meters are having no effect at all on whether and to what extent our creature believes it to be less than twenty meters. So in what way *are* the beliefs that the distance is or appears to be less than three meters serving for the creature as a basis for his belief that the distance is less than twenty meters?

One way for a strange all-or-nothing believer to base his belief that the distance is less than twenty meters on something like the fact that it appears to be less than three meters — and also to start looking like someone who knows the distance is less than twenty meters—is for him to have learned the trick of not believing (at all) such things as that a distance is less than twenty meters unless he's inclined to believe that, or unless it appears to him that, the distance is considerably less than that as well. But, again, this very change makes him cease to be the kind of

²² In connection with this, see (DeRose 2005: 160), where I take a similar approach to the closely related matter of when one belief transmits warrant to another.

²³ But then we face the issue of what we *will* require in terms of how Q supports P, which, as I remarked above in note 17, is a very tricky matter.

distance underestimator who causes trouble for insensitivity accounts in the first place. If we keep him a distance underestimator of the relevant type, and thereby imagine that he's not come to practice the trick just described, and we deprive him of the ability to have the degree to which he believes that the distance is less than twenty meters to be at all affected in the degree to which it's held by the fact that it also appears to be considerably less than that, it becomes difficult to see how he is in any good way basing his belief that the distance is less than twenty meters on its giving him appearances of being considerably less than that.

So, my own suspicion is that RI will in the end, like the degree-of-belief insensitivity account, predict that this strange creature will seem not to know that the distance he sees is less than twenty meters. But that's alright, since the very considerations that make the creature run afoul of these accounts are also considerations that push strongly in the direction of ruling that the creature doesn't know that the distance is less than twenty meters — which I think would be the intuition about this case for those who fully come to grips with what such a creature is like.

But it may well happen that those who see matters very differently from me here and are unwavering about the intuition being that the all-or-nothing-believing, underestimating creature knows that the distance he is seeing is less than twenty meters will also be inclined to think that this creature can very well be basing his belief that the distance is less than twenty meters on such beliefs as that it appears to him to be less than three meters, despite the creature's severe cognitive deficiencies. Though for quite different reasons from mine, these folks, too, then, will find RI to survive this example.²⁴

²⁴ For discussion and comments, thanks to the participants in the seminar on 'Safety and Sensitivity of Beliefs' that Troy Cross and I co-taught at Yale in the Spring of 2006, and especially to Troy and to Geoff Pynn. Thanks also to the audience at my presentation at the conference on Saul Kripke's *Philosophical Troubles* at the Saul Kripke Center, CUNY Graduate Center; 15 September 2011.